

Faster SVD-Truncated Regularized Least-Squares

Christos Boutsidis*

Malik Magdon-Ismail[†]

Abstract

We develop a fast algorithm for computing the “SVD-truncated” regularized solution to the least-squares problem: $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$. Let \mathbf{A}_k of rank k be the best rank k matrix computed via the SVD of \mathbf{A} . Then, the SVD-truncated regularized solution is: $\mathbf{x}_k = \mathbf{A}_k^\dagger \mathbf{b}$. If \mathbf{A} is $m \times n$, then, it takes $O(mn \min\{m, n\})$ time to compute \mathbf{x}_k using the SVD of \mathbf{A} . We give an approximation algorithm for \mathbf{x}_k which constructs a rank- k approximation $\tilde{\mathbf{A}}_k$ and computes $\tilde{\mathbf{x}}_k = \tilde{\mathbf{A}}_k^\dagger \mathbf{b}$ in roughly $O(\text{nnz}(\mathbf{A})k \log n)$ time. Our algorithm uses a randomized variant of the subspace iteration. We show that, with high probability: $\|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|_2 \approx \|\mathbf{Ax}_k - \mathbf{b}\|_2$ and $\|\mathbf{x}_k - \tilde{\mathbf{x}}_k\|_2 \approx 0$.

1 Introduction

We consider the least-squares regression problem:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{Ax} - \mathbf{b}\|_2, \quad \mathbf{A} \in \mathbb{R}^{m \times n}, \quad \mathbf{b} \in \mathbb{R}^m, \quad m \geq n.$$

We assume \mathbf{A} has rank $\rho \leq n$. Via the SVD (see Section 3), we can decompose \mathbf{A} as:

$$\mathbf{A} = \sum_{i=1}^{\rho} \sigma_i \mathbf{u}_i \mathbf{v}_i^T,$$

where $\{\mathbf{u}_i \in \mathbb{R}^m, \mathbf{v}_i \in \mathbb{R}^n, \sigma_i \in \mathbb{R}_+\}$ are respectively the left and right singular vectors, and singular values of \mathbf{A} . The orthonormal left and right singular matrices $\mathbf{U}_\mathbf{A} = [\mathbf{u}_1, \dots, \mathbf{u}_\rho]$ and $\mathbf{V}_\mathbf{A} = [\mathbf{v}_1, \dots, \mathbf{v}_\rho]$ have the singular vectors as columns. The optimal solution to this least-squares problem is:

$$\mathbf{x}_{opt} = \sum_{i=1}^{\rho} \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i.$$

When $\sigma_i \rightarrow 0$, for some i , the solution is numerically unstable, and the problem becomes ill-posed [10]. Regularization helps with the numerical instability as well as improving the generalization performance of machine learning algorithms that use regression. Perhaps the simplest and most popular regularization technique is Tikhonov regularization (or weight decay) [1, Ex. 4.5] which results in the solution:

$$\mathbf{x}_\lambda = \sum_{i=1}^{\rho} \frac{\sigma_i^2}{\sigma_i^2 + \lambda_i^2} \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i,$$

where $\lambda_i > 0$ are the regularization parameters (often chosen uniform, $\lambda_i = \lambda$). This regularized solution minimizes a penalized error $\|\mathbf{Ax} - \mathbf{b}\|_2^2 + \|\mathbf{A}\mathbf{V}_\mathbf{A}^T \mathbf{x}\|_2^2$, where $\mathbf{\Lambda}$ is a $\rho \times \rho$ diagonal matrix with entries $\mathbf{\Lambda}_{ii} = \lambda_i$; each λ_i quantifies how much one chooses to regularize the i th singular space of \mathbf{A} .

*Mathematical Sciences Department, IBM T.J. Watson Research Center, Yorktown Heights NY 10598, USA.

[†]Computer Science Department, Rensselaer Polytechnic Institute, Troy NY 12180, USA, magdon@cs.rpi.edu.

SVD-truncated regularization, the focus of this paper, is a special case of Tikhonov regularization with $\lambda_i = 0$ for $i \leq k$ ($k < \rho$ is the truncation parameter) and $\lambda_i \rightarrow \infty$ otherwise [10]. The SVD-truncated regularized solution \mathbf{x}_k is:

$$\mathbf{x}_k = \sum_{i=1}^k \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i.$$

From the SVD, $\mathbf{A}_k = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^T$ is a best rank k approximation to \mathbf{A} . So, \mathbf{x}_k solves a least-squares problem with \mathbf{A}_k :

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{A}_k \mathbf{x} - \mathbf{b}\|_2, \quad \mathbf{A} \in \mathbb{R}^{m \times n}, \quad \mathbf{b} \in \mathbb{R}^m, \quad m \geq n.$$

Appropriately choosing λ_i or k are important problems from the numerical linear algebra perspective as well as the machine learning perspective, and we refer to Section 5 in [10] for some discussion on this topic. For our purposes, we take k as given. That is, \mathbf{x}_k is the solution we want¹, and our goal is to compute a good approximation to \mathbf{x}_k quickly in $o(mn^2)$ time, since the SVD may be too expensive if \mathbf{A} is massive.

Our contributions. Via a recent randomized variant [13, 8] of the subspace iteration method, we develop a fast randomized algorithm to compute an $\tilde{\mathbf{x}}_k$ in roughly $O(\text{nnz}(\mathbf{A})k \log n)$ time where $\text{nnz}(\mathbf{A})$ is the number of non-zeros in \mathbf{A} . We describe this algorithm in Section 4 and give precise error estimates for its performance in Theorem 5. We show that there is not much room for improvement upon these estimates by providing a lower bound in Theorem 9.

2 Related Work

SVD-truncated regression has been around for some time. See, for example, [16, 10] and references therein for some background and applications of this regularization technique.

To develop faster SVD-truncated regression, our approach is to first compute $\tilde{\mathbf{A}}_k$, an approximation to \mathbf{A}_k , obviously to \mathbf{b} , and use $\tilde{\mathbf{A}}_k$ in the regression. To construct $\tilde{\mathbf{A}}_k$, we use an algorithm that was previously proposed to quickly construct a “good” low-rank approximation to a matrix in the spectral norm. This algorithm is based on the subspace iteration method [6, Sec 8.2.4] and was analyzed in [13], [8].

The approach of first approximating \mathbf{A}_k is natural and has been used before, for example [18] for uniform Tikhonov regularization ($\lambda_i = \lambda$). The main algorithm in [18] is similar to ours without the power iteration, and corresponds to a random embedding into $k + q$ dimensions before computing an approximate basis for the column space of \mathbf{A} . Theorem 1 in [18] provides a high probability bound:

$$\|\mathbf{x}_\lambda - \hat{\mathbf{x}}_\lambda\|_2^2 = O((k + q)(q \log q + n - k)\gamma_k^2) \cdot \|\mathbf{x}_\lambda\|_2^2.$$

Here, $\gamma_k = \sigma_{k+1}(\mathbf{A})/\sigma_k(\mathbf{A}) \leq 1$. This bound is similar in spirit to our Eqn. (2) in Theorem 5, except we work with SVD-truncated regularization, not uniform Tikhonov regularization, and we give a stronger $O(\varepsilon)$ bound.

The approach in [17] also uses subspace iteration as we do, with a different choice for the dimension reduction and an orthonormalization step (see Section 3.1 in [17]) - this choice is the “classical subspace iteration method” from the numerical linear algebra literature [6, Sec 8.2.4]. However, no theoretical bounds are reported in [17]. Iterative SVD-based methods such as the Lanczos iteration were also proposed in Section 4 in [17] and [12]. These approaches enjoy good empirical behavior, but again, no theoretical bounds are known.

¹When \mathbf{b} is concentrated in the top singular subspaces, \mathbf{x}_k approximates \mathbf{x}_{opt} . Indeed, if $\mathbf{u}_i^T \mathbf{b} \geq \sigma_i^\alpha$ for some $\alpha \geq 1$, then ([10, Theorem 3.1]): $\|\mathbf{x}_{opt} - \mathbf{x}_k\|_2 / \|\mathbf{x}_{opt}\|_2 \leq \sqrt{n} (\sigma_{k+1}/\sigma_k)^{\alpha-1}$. That is, the SVD-truncated solution is near optimal when the singular value gap $\gamma_k = \sigma_{k+1}/\sigma_k$ is small.

In the above two results [17, 12], when we say there are no theoretical bounds we mean there are no bounds for the regression setting, as those we provide in Theorem 5. However, subspace iteration and Lanczos iteration were extensively analyzed before and bounds similar to Lemma 7 are available.

An alternative approach to SVD-truncation is feature selection or sparsity. In this setting, one selects columns from \mathbf{A} and solves the reduced regression with only these columns, resulting in a sparse solution. See Section 12.2 in [6] for a discussion of this approach. In recent work [2], we developed a method based on column sampling that runs in $O(mn \min\{m, n\} + nk^3/\varepsilon^2)$ time and returns a solution $\hat{\mathbf{x}}_r \in \mathbb{R}^n$ with $r = O(k)$ non-zero entries such that:

$$\|\mathbf{A}\hat{\mathbf{x}}_r - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2 + O(1)\|\mathbf{b}\|_2\|\mathbf{A} - \mathbf{A}_k\|_F\sigma_k^{-1}(\mathbf{A}).$$

Eqn. 1 in Theorem 5 in the present article, when we remove the sparsity constraint, is considerably tighter.

A similar bound can be obtained using the Rank-Revealing QR (RRQR) factorization [3]: a QR-like decomposition is used to select exactly k columns of \mathbf{A} to obtain a sparse solution $\hat{\mathbf{x}}_k$. Combining Eqn. (12) of [3] with Strong RRQR [7] gives

$$\|\mathbf{x}_k - \hat{\mathbf{x}}_k\|_2 \leq 3 \left(\sqrt{4k(n-k)} + 1 \right) \sigma_k^{-1}(\mathbf{A}) \cdot \|\mathbf{b}\|_2.$$

Eqn. 2 in Theorem 5 in the present article, when we remove the sparsity constraint, is considerably tighter.

3 Preliminaries

Basic Notation. We use $\mathbf{A}, \mathbf{B}, \dots$ to denote matrices; $\mathbf{a}, \mathbf{b}, \dots$ to denote column vectors. \mathbf{I}_n is the $n \times n$ identity matrix; $\mathbf{0}_{m \times n}$ is the $m \times n$ matrix of zeros. We use $\|\mathbf{A}\|_F$ for the Frobenius matrix norm and $\|\mathbf{A}\|_2$ for the spectral or operator norm: $\|\mathbf{A}\|_F^2 = \sum_{i,j} \mathbf{A}_{ij}^2$ and $\|\mathbf{A}\|_2 = \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$. By submultiplicativity, $\|\mathbf{A}\mathbf{B}\|_2 \leq \|\mathbf{A}\|_2\|\mathbf{B}\|_2$, for any \mathbf{A}, \mathbf{B} .

Singular Value Decomposition and the Pseudo-inverse. The thin (compact) Singular Value Decomposition (SVD) of the matrix \mathbf{A} with $\text{rank}(\mathbf{A}) = \rho$ is:

$$\mathbf{A} = \underbrace{\begin{pmatrix} \mathbf{U}_k & \mathbf{U}_{\rho-k} \end{pmatrix}}_{\mathbf{U}_\mathbf{A} \in \mathbb{R}^{m \times \rho}} \underbrace{\begin{pmatrix} \Sigma_k & \mathbf{0} \\ \mathbf{0} & \Sigma_{\rho-k} \end{pmatrix}}_{\Sigma_\mathbf{A} \in \mathbb{R}^{\rho \times \rho}} \underbrace{\begin{pmatrix} \mathbf{V}_k^\top \\ \mathbf{V}_{\rho-k}^\top \end{pmatrix}}_{\mathbf{V}_\mathbf{A}^\top \in \mathbb{R}^{\rho \times n}},$$

where $\Sigma_\mathbf{A}$, a positive diagonal matrix, contains the singular values in decreasing order: $(\Sigma_\mathbf{A})_{ii} = \sigma_i(\mathbf{A})$ (we will drop the dependence on \mathbf{A} and use σ_i when the context is clear). The matrices $\mathbf{U}_k \in \mathbb{R}^{m \times k}$ and $\mathbf{U}_{\rho-k} \in \mathbb{R}^{m \times (\rho-k)}$ contain the left singular vectors of \mathbf{A} ; similarly, $\mathbf{V}_k \in \mathbb{R}^{n \times k}$ and $\mathbf{V}_{\rho-k} \in \mathbb{R}^{n \times (\rho-k)}$ contain the right singular vectors of \mathbf{A} . $\mathbf{A}_k = \mathbf{U}_k \Sigma_k \mathbf{V}_k^\top = \mathbf{U}_k \mathbf{U}_k^\top \mathbf{A} = \mathbf{A} \mathbf{V}_k \mathbf{V}_k^\top \in \mathbb{R}^{m \times n}$ minimizes $\|\mathbf{A} - \mathbf{X}\|_2$ over all matrices $\mathbf{X} \in \mathbb{R}^{m \times n}$ of rank at most k . Note that $\|\mathbf{A}\|_2 = \sigma_1(\mathbf{A})$ and $\|\mathbf{A} - \mathbf{A}_k\|_2 = \sigma_{k+1}(\mathbf{A})$. The pseudo-inverse of \mathbf{A} is $\mathbf{A}^\dagger = \mathbf{V}_\mathbf{A} \Sigma_\mathbf{A}^{-1} \mathbf{U}_\mathbf{A}^\top \in \mathbb{R}^{n \times m}$. The spectral gap of \mathbf{A} at $k < \text{rank}(\mathbf{A})$ is $\gamma_k = \sigma_{k+1}(\mathbf{A})/\sigma_k(\mathbf{A}) \leq 1$.

Perturbation Theory. There exist bounds on the perturbation of the pseudoinverse and singular values of a matrix upon additive perturbation. Let $\mathbf{A}, \mathbf{B}, \mathbf{E}$ be $m \times n$ matrices with $\mathbf{B} = \mathbf{A} + \mathbf{E}$.

Lemma 1 ([15, Theorem 3.4]). *If $m \geq n$ and $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{B}) < \min\{m, n\}$: $\|\mathbf{B}^\dagger - \mathbf{A}^\dagger\|_2 \leq 2\|\mathbf{A}^\dagger\|_2\|\mathbf{B}^\dagger\|_2\|\mathbf{E}\|_2$.*

Lemma 2 (Weyl's inequality [11, Corollary 7.3.8]). *$|\sigma_i(\mathbf{B}) - \sigma_i(\mathbf{A})| \leq \|\mathbf{E}\|_2$, for $i = 1, 2, \dots, \min(m, n)$.*

Random Matrix Theory. There exist results bounding the top and bottom singular values of a random Gaussian matrix.

Lemma 3 (Norm of a Gaussian Matrix [4]). *Let $\mathbf{X} \in \mathbb{R}^{n \times m}$ be a matrix of i.i.d. standard Gaussian random variables, where $n \geq m$. Then, for $t \geq 4$, $\mathbb{P}\{\sigma_1(\mathbf{X}) \geq tn^{\frac{1}{2}}\} \geq e^{-nt^2/8}$.*

Lemma 4 (Invertibility of a Gaussian Matrix [14]). *Let $\mathbf{X} \in \mathbb{R}^{n \times n}$ be a matrix with i.i.d. standard Gaussian random variables. Then, for $\delta > 0$, $\mathbb{P}\{\sigma_n(\mathbf{X}) \leq \delta n^{-\frac{1}{2}}\} \leq 2.35\delta$.*

4 Main Result

We use an approximation $\tilde{\mathbf{A}}_k$ (instead of \mathbf{A}_k) and minimize $\|\tilde{\mathbf{A}}_k \mathbf{x} - \mathbf{b}\|_2$ over \mathbf{x} . The algorithm is summarized below.

1. Compute $\mathbf{Q} \in \mathbb{R}^{m \times k}$, an orthonormal basis for the columns of $(\mathbf{A}\mathbf{A}^T)^p \mathbf{A}\mathbf{S} \in \mathbb{R}^{m \times k}$, where $p \geq 0$ and \mathbf{S} is an $n \times k$ matrix of i.i.d. standard Gaussians.
2. Compute $\tilde{\mathbf{A}}_k = \mathbf{Q}\mathbf{Q}^T \mathbf{A}$ and $\tilde{\mathbf{x}}_k = \tilde{\mathbf{A}}_k^\dagger \mathbf{b}$.

Careful implementation makes the algorithm efficient. In Step 1, we compute the matrix products in $(\mathbf{A}\mathbf{A}^T)^p \mathbf{A}\mathbf{S}$ from right to left to ensure a running time of $O(\text{nnz}(\mathbf{A})kp)$; the result is an $m \times k$ matrix, and a QR-factorization in time $O(mk^2)$ gives \mathbf{Q} . In Step 2, we need the SVD $\tilde{\mathbf{A}}_k = \tilde{\mathbf{U}}_k \tilde{\Sigma}_k \tilde{\mathbf{V}}_k^T$. Instead, we compute the SVD of $\mathbf{Q}^T \mathbf{A} = \mathbf{U}_{\mathbf{Q}^T \mathbf{A}} \Sigma_{\mathbf{Q}^T \mathbf{A}} \mathbf{V}_{\mathbf{Q}^T \mathbf{A}}^T$, in $O(\text{nnz}(\mathbf{A})k + nk^2)$ time². Then, $\tilde{\mathbf{A}}_k = \mathbf{Q}\mathbf{U}_{\mathbf{Q}^T \mathbf{A}} \Sigma_{\mathbf{Q}^T \mathbf{A}} \mathbf{V}_{\mathbf{Q}^T \mathbf{A}}^T$, from which we read off the SVD of $\tilde{\mathbf{A}}_k$ because $\mathbf{Q}\mathbf{U}_{\mathbf{Q}^T \mathbf{A}}$ is orthonormal: $\tilde{\mathbf{U}}_k = \mathbf{Q}\mathbf{U}_{\mathbf{Q}^T \mathbf{A}}$, $\tilde{\Sigma}_k = \Sigma_{\mathbf{Q}^T \mathbf{A}}$, $\tilde{\mathbf{V}}_k = \mathbf{V}_{\mathbf{Q}^T \mathbf{A}}$. Now, $\tilde{\mathbf{A}}_k^\dagger = \tilde{\mathbf{V}}_k \tilde{\Sigma}_k^{-1} \tilde{\mathbf{U}}_k^T$ and $\tilde{\mathbf{x}}_k = \tilde{\mathbf{V}}_k \tilde{\Sigma}_k^{-1} \tilde{\mathbf{U}}_k^T \mathbf{b}$, which is computed in $O(mk + nk + k^2)$ time. The dominant terms in the running time are $O(\text{nnz}(\mathbf{A})kp + (m+n)k^2)$.

We control the accuracy of the algorithm by choosing p appropriately. A larger p gives a better error. Let $0 < \varepsilon < 1$ be an error parameter and recall that the spectral gap of \mathbf{A} at k is $\gamma_k = \sigma_{k+1}(\mathbf{A})/\sigma_k(\mathbf{A}) \leq 1$. The next theorem quantifies how the error depends on p . Roughly speaking, setting $p = O(\ln(\varepsilon/n)/\ln(\gamma_k))$ suffices to give additive error $\varepsilon\|\mathbf{b}\|_2$.

Theorem 5. *Fix $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, $k < \text{rank}(\mathbf{A})$, and $\varepsilon, \delta \in (0, 1)$. Choose p in our algorithm to satisfy*

$$p \geq \frac{\ln(\varepsilon \cdot \delta \cdot \frac{\sigma_k^2}{\sigma_1^2} \cdot \frac{1}{12n})}{\ln(\gamma_k^2)}.$$

Let $\tilde{\mathbf{x}}_k = \tilde{\mathbf{A}}_k^\dagger \mathbf{b}$ and $\mathbf{x}_k = \mathbf{A}_k^\dagger \mathbf{b}$ be the exact SVD-truncated solution. Then, with probability at least $1 - e^{-2n} - 2.35\delta$:

$$\|\mathbf{A}\tilde{\mathbf{x}}_k - \mathbf{b}\|_2 \leq \|\mathbf{A}\mathbf{x}_k - \mathbf{b}\|_2 + \varepsilon \cdot \|\mathbf{b}\|_2, \quad (1)$$

and

$$\frac{\|\mathbf{x}_k - \tilde{\mathbf{x}}_k\|_2}{\|\mathbf{x}_k\|_2} \leq \frac{4}{3} \cdot \varepsilon. \quad (2)$$

The error in (1) is additive, and in Section 5 we show that this is unavoidable when, as we do, one solves the regression via an approximation $\tilde{\mathbf{A}}_k$ which is constructed obliviously to \mathbf{b} .

²It takes $O(\text{nnz}(\mathbf{A})k)$ time for the matrix multiplication $\mathbf{Q}^T \mathbf{A}$, and $O(nk^2)$ time to compute the SVD of $\mathbf{Q}^T \mathbf{A}$.

4.1 Proof of Theorem 5

Recall that $\mathbf{A}_k = \mathbf{U}_k \mathbf{U}_k^T \mathbf{A}$ and $\mathbf{x}_k = \mathbf{A}_k^\dagger \mathbf{b}$. By our construction of $\tilde{\mathbf{A}}_k$ and $\tilde{\mathbf{U}}_k$, $\tilde{\mathbf{A}}_k = \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T \mathbf{A}$, with $\tilde{\mathbf{x}}_k = \tilde{\mathbf{A}}_k^\dagger \mathbf{b}$. We first quantify the additive error. By the triangle inequality,

$$\|\mathbf{A} \tilde{\mathbf{x}}_k - \mathbf{b}\|_2 \leq \|\mathbf{A} \mathbf{x}_k - \mathbf{b}\|_2 + \Delta,$$

where

$$\Delta = \|\mathbf{A} \tilde{\mathbf{x}}_k - \mathbf{A} \mathbf{x}_k\|_2 = \|\mathbf{A}(\tilde{\mathbf{A}}_k^\dagger - \mathbf{A}_k^\dagger) \mathbf{b}\|_2.$$

We need to upper bound Δ . By submultiplicativity,

$$\begin{aligned} \|\tilde{\mathbf{A}}_k - \mathbf{A}_k\|_2 &= \|(\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T) \mathbf{A}\|_2 \\ &\leq \|\mathbf{A}\|_2 \|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2. \end{aligned} \quad (3)$$

Lemma 6. $\Delta \leq \frac{2\sigma_1^2(\mathbf{A})}{\sigma_k(\tilde{\mathbf{A}})\sigma_k(\mathbf{A})} \|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2 \|\mathbf{b}\|_2$.

Proof. We manipulated Δ as follows:

$$\begin{aligned} \Delta &= \|\mathbf{A}(\tilde{\mathbf{A}}_k^\dagger - \mathbf{A}_k^\dagger) \mathbf{b}\|_2 \\ &\leq \|\mathbf{A}\|_2 \|\tilde{\mathbf{A}}_k^\dagger - \mathbf{A}_k^\dagger\|_2 \|\mathbf{b}\|_2 \\ &\leq 2\|\mathbf{A}\|_2 \|\tilde{\mathbf{A}}_k^\dagger\|_2 \|\mathbf{A}_k^\dagger\|_2 \|\tilde{\mathbf{A}}_k - \mathbf{A}_k\|_2 \|\mathbf{b}\|_2 \\ &\leq \frac{2\sigma_1^2(\mathbf{A})}{\sigma_k(\tilde{\mathbf{A}})\sigma_k(\mathbf{A})} \|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2 \|\mathbf{b}\|_2 \end{aligned}$$

The first inequality uses submultiplicativity; the second uses Lemma 1; and, the last uses Eqn. (3). ■

Lemma 6 holds no matter what $\tilde{\mathbf{U}}_k$ is. The difference in the projection operators $\|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2$ plays an important role in our bounds. Our algorithm constructs $\tilde{\mathbf{U}}_k$ for which this error term can be bounded. A similar application of the power iteration was analyzed for spectral clustering in [5]. For the specific $\tilde{\mathbf{U}}_k$ returned by our algorithm, the difference in the projection operators can be bounded with high probability.

Lemma 7. [5] Fix $\varepsilon, \delta \in (0, 1)$. If $p \geq \ln(\varepsilon\delta/4n)/\ln(\gamma_k^2)$, then with probability at least $1 - e^{-2n} - 2.35\delta$,

$$\|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2 \leq \varepsilon.$$

Proof. See the Appendix. ■

Lemma 7 bounds the difference in the projection operators. Notice that in Lemma 6 we also need the k th singular value of $\tilde{\mathbf{A}}_k$. This can be bounded by Weyl's theorem (Lemma 2):

$$\begin{aligned} |\sigma_k(\tilde{\mathbf{A}}_k) - \sigma_k(\mathbf{A}_k)| &\leq \|\tilde{\mathbf{A}}_k - \mathbf{A}_k\|_2 \\ &\leq \sigma_1(\mathbf{A}) \|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2, \end{aligned}$$

from which we have that

$$\sigma_k(\tilde{\mathbf{A}}_k) \geq \sigma_k - \sigma_1 \|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2. \quad (4)$$

We are now ready to prove Eqn. (1) in Theorem 5. Set

$$p \geq \frac{\ln(\varepsilon \delta \sigma_k^2 / 12 n \sigma_1^2)}{\ln(\gamma_k^2)} = \frac{\ln\left(\left(\frac{\varepsilon \sigma_k^2}{3 \sigma_1^2}\right) \delta / 4 n\right)}{\ln(\gamma_k^2)}.$$

It now follows from Lemma 7 that

$$\|\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T - \mathbf{U}_k \mathbf{U}_k^T\|_2 \leq \frac{\varepsilon \sigma_k^2}{3 \sigma_1^2}. \quad (5)$$

From (4), it follows that

$$\sigma_k(\tilde{\mathbf{A}}_k) \geq \sigma_k - \frac{\varepsilon \sigma_k^2}{3 \sigma_1^2} = \sigma_k \left(1 - \frac{\varepsilon \sigma_k}{3 \sigma_1}\right) \geq \frac{2}{3} \sigma_k. \quad (6)$$

Using Lemma 6 with (5) and (6) we obtain a bound for Δ :

$$\Delta \leq \frac{2 \sigma_1^2}{\frac{2}{3} \sigma_k \cdot \sigma_k} \cdot \left(\frac{\varepsilon \sigma_k^2}{3 \sigma_1^2}\right) \cdot \|\mathbf{b}\|_2 = \varepsilon \|\mathbf{b}\|_2.$$

We move to the proof of Eqn. (2) in Theorem 5. We need a perturbation theory result from [9] which we state in our notation for a perturbation of the matrix \mathbf{A}_k to the matrix $\tilde{\mathbf{A}}_k$, without any perturbation on the response \mathbf{b} . Let

$$\mathbf{E} = \tilde{\mathbf{A}}_k - \mathbf{A}_k.$$

Lemma 8 ([9, Eqn. (27)]). *If $\|\mathbf{E}\|_2 < \sigma_k$, then,*

$$\frac{\|\mathbf{x}_k - \tilde{\mathbf{x}}_k\|_2}{\|\mathbf{x}_k\|_2} \leq \frac{\sigma_1 \|\mathbf{E}\|_2}{\sigma_k - \|\mathbf{E}\|_2} \left(\frac{1}{\sigma_1} + \frac{\|\mathbf{A}_k \mathbf{x}_k - \mathbf{b}\|_2}{\sigma_k \|\mathbf{b}\|_2} \right) + \frac{\|\mathbf{E}\|_2}{\sigma_k}.$$

We can simplify the bound in Lemma 8 because $\|\mathbf{A}_k \mathbf{x}_k - \mathbf{b}\|_2 \leq \|\mathbf{b}\|_2$:

$$\frac{\|\mathbf{x}_k - \tilde{\mathbf{x}}_k\|_2}{\|\mathbf{x}_k\|_2} \leq \frac{\sigma_1 \|\mathbf{E}\|_2}{\sigma_k - \|\mathbf{E}\|_2} \left(\frac{1}{\sigma_1} + \frac{1}{\sigma_k} \right) + \frac{\|\mathbf{E}\|_2}{\sigma_k}. \quad (7)$$

Using the bound (5) in (3) and recalling that $\gamma_k = \frac{\sigma_{k+1}}{\sigma_k}$,

$$\|\mathbf{E}\| = \|\tilde{\mathbf{A}}_k - \mathbf{A}_k\|_2 \leq \frac{\varepsilon \gamma_k}{3} \cdot \sigma_k.$$

In particular, since $\varepsilon < 1$ and $\gamma_k \leq 1$, $\|\tilde{\mathbf{A}}_k - \mathbf{A}_k\|_2 < \sigma_k$ and we can apply Lemma 8, or the bound in Eqn. (7):

$$\begin{aligned} \frac{\|\mathbf{x}_k - \tilde{\mathbf{x}}_k\|_2}{\|\mathbf{x}_k\|_2} &\leq \frac{\sigma_1 \left(\frac{\varepsilon \gamma_k}{3} \sigma_k\right)}{\sigma_k - \left(\frac{\varepsilon \gamma_k}{3} \sigma_k\right)} \left(\frac{1}{\sigma_1} + \frac{1}{\sigma_k} \right) + \frac{\left(\frac{\varepsilon \gamma_k}{3} \sigma_k\right)}{\sigma_k} \\ &= \frac{\varepsilon \gamma_k}{3} \cdot \left(\frac{1}{1 - \frac{\varepsilon \gamma_k}{3}} \left(1 + \frac{1}{\gamma_k} \right) + 1 \right) \\ &\leq \frac{\varepsilon \gamma_k}{3} \cdot \left(\frac{3}{2} \left(1 + \frac{1}{\gamma_k} \right) + 1 \right) \\ &= \frac{\varepsilon}{3} \cdot \left(\frac{5 \gamma_k}{2} + \frac{3}{2} \right) \leq \frac{4}{3} \varepsilon. \end{aligned}$$

(The second inequality is because $1 - \frac{\varepsilon \gamma_k}{3} \geq \frac{2}{3}$; and, the final inequality is because $\gamma_k \leq 1$.) ■

5 Lower Bound: Additive Error is Unavoidable

We now show that the additive error of Eqn. 1 in Theorem 5 is tight. Towards this end, let us consider the class of (fast) algorithms which operate as follows:

1. Quickly construct matrix $\tilde{\mathbf{A}}_k$ of rank k obliviously to \mathbf{b} .
2. Use $\tilde{\mathbf{A}}_k$ to construct the approximate solution $\tilde{\mathbf{x}}_k = \tilde{\mathbf{A}}_k^\dagger \mathbf{b}$.

Let $\|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k\|_2 = \varepsilon \|\mathbf{A}_k\|_2$. The cross-‘projection’ operator $\mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger$ quantifies how well $\tilde{\mathbf{A}}_k^\dagger$ approximates \mathbf{A}_k^\dagger (if $\tilde{\mathbf{A}}_k^\dagger = \mathbf{A}_k^\dagger$, then it is a projection operator and $\varepsilon = 0$). Note that $\tilde{\mathbf{A}}_k^\dagger - \mathbf{A}_k^\dagger$ and $\mathbf{A}_k - \tilde{\mathbf{A}}_k = (\mathbf{U}_k \mathbf{U}_k^\top - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^\top) \mathbf{A}$ are related (see the discussion in Section 1), and for our algorithm $\|\mathbf{U}_k \mathbf{U}_k^\top - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^\top\|_2$ is bounded by ε (see Lemma 7).

The next theorem states that the additive error in Eqn. 1 in Theorem 5 is about the best you can expect of algorithms that construct $\tilde{\mathbf{x}}_k$ via an approximation $\tilde{\mathbf{A}}_k$, provided that $\tilde{\mathbf{A}}_k$ is constructed obliviously to \mathbf{b} . The notion of approximation we consider is via the equation $\|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k\|_2 = \varepsilon \|\mathbf{A}_k\|_2$.

Theorem 9. Fix $\mathbf{A}, \tilde{\mathbf{A}}_k \in \mathbb{R}^{m \times n}$. Let \mathbf{A}_k be the best rank- k approximation to \mathbf{A} used in the top- k SVD-truncated regression and suppose $\tilde{\mathbf{A}}_k$ satisfies (for $\varepsilon > 0$):

$$\|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k\|_2 = \varepsilon \|\mathbf{A}_k\|_2.$$

Then, for some $\mathbf{b} \in \mathbb{R}^m$, with $\mathbf{x}_k = \mathbf{A}_k^\dagger \mathbf{b}$ and $\tilde{\mathbf{x}}_k = \tilde{\mathbf{A}}_k^\dagger \mathbf{b}$,

$$\begin{aligned} \|\mathbf{A} \mathbf{x}_k - \mathbf{b}\|_2 &= 0 \\ \|\mathbf{A} \tilde{\mathbf{x}}_k - \mathbf{b}\|_2 &\geq \varepsilon \|\mathbf{b}\|_2. \end{aligned}$$

In particular, no multiplicative error bound is possible and the additive error is at least $\varepsilon \|\mathbf{b}\|_2$.

Proof. We set $\mathbf{b} = \mathbf{A}_k \mathbf{z}$ for \mathbf{z} to be selected later. Then

$$\mathbf{A} \mathbf{x}_k - \mathbf{b} = (\mathbf{A} \mathbf{A}_k^\dagger \mathbf{A}_k - \mathbf{A}_k) \mathbf{z} = \mathbf{0}.$$

(The last equality is because $\mathbf{A} \mathbf{A}_k^\dagger \mathbf{A}_k = \mathbf{A}_k \mathbf{A}_k^\dagger \mathbf{A}_k = \mathbf{A}_k$.) We now manipulate $\|\mathbf{A} \tilde{\mathbf{x}}_k - \mathbf{b}\|_2$.

$$\begin{aligned} \|\mathbf{A} \tilde{\mathbf{x}}_k - \mathbf{b}\|_2 &= \|\mathbf{A} \tilde{\mathbf{A}}_k^\dagger \mathbf{A}_k \mathbf{z} - \mathbf{A}_k \mathbf{z}\|_2 \\ &= \|\mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger \mathbf{A}_k \mathbf{z} - \mathbf{A}_k \mathbf{z} + \mathbf{A}_{\rho-k} \tilde{\mathbf{A}}_k^\dagger \mathbf{A}_k \mathbf{z}\|_2 \\ &\geq \|\mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger \mathbf{A}_k \mathbf{z} - \mathbf{A}_k \mathbf{z}\|_2 \\ &= \|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k \mathbf{z}\|_2 \end{aligned}$$

(The inequality follows because $\|\mathbf{X} + \mathbf{Y}\|_2 \geq \|\mathbf{X}\|_2$ when $\mathbf{X}^\top \mathbf{Y} = \mathbf{0}$.) We now choose \mathbf{z} to be the top right singular vector of the matrix $(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k$. Then,

$$\begin{aligned} \|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k \mathbf{z}\|_2 &= \|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k\|_2 \cdot \|\mathbf{z}\|_2 \\ &\geq \|(\mathbf{I} - \mathbf{A}_k \tilde{\mathbf{A}}_k^\dagger) \mathbf{A}_k\|_2 \cdot \frac{\|\mathbf{b}\|_2}{\|\mathbf{A}_k\|_2} \\ &= \varepsilon \|\mathbf{b}\|_2. \end{aligned}$$

The inequality uses $\|\mathbf{b}\|_2 = \|\mathbf{A}_k \mathbf{z}\|_2 \leq \|\mathbf{A}_k\|_2 \|\mathbf{z}\|_2$, and the last equality is from the theorem statement. ■




Figure 1: (a) Solid curves: accuracy ratio for our fast SVD-truncated regression algorithm, $k = 20$, $p = 20 \ln n$. Dashed curves (for comparison): MATLAB's `svds` solver which computes an approximation to the top- k singular space; we set the error tolerance to produce comparable error to our algorithm (options.tol=0.001). The performance of our algorithm results in nearly constant error (approx 4% in the objective and 1% in the solution vector); the accuracy of `svds` is approximately the same but very unpredictable. (b) Ratio of the time to compute the exact solution over time to compute the approximation. For $n = 1500$ our algorithm is about 50% more efficient. MATLAB's `svds` is also more efficient asymptotically than the exact solution, but is not as efficient as our algorithm.

6 Numerical Illustration

We perform a numerical experiment on a synthetic regression problem to illustrate the theory and the algorithm. We construct a synthetic problem as follows. We generate an $n \times n$ matrix A of i.i.d. Gaussians, and set the spectral gap $\gamma_k = \sigma_{k+1}/\sigma_k = 0.99$. To do this, use the SVD, $A = U\Sigma V^T$ and rescale $\sigma_{k+1}, \dots, \sigma_n$ up or down by a constant factor so that $\gamma_k = 0.99$. Now reconstruct A using U , V and the rescaled Σ . We construct the response $b = \frac{A_k r_1}{\|A_k r_1\|_2} + 0.2 \times \frac{r_2}{\|r_2\|_2}$, where r_1, r_2 are random standard Gaussian vectors. So, the response b has roughly 80% within the top- k singular space. We set $k = 20$ and run our algorithm with $p = 10 \ln n$. We vary $n \in [100, 1000]$ and for each value of n take the average over several experiments to increase statistical significance.

For comparison, we also use the truncated SVD algorithm `svds` distributed with MATLAB 8.1, where one can specify an error tolerance `tol`; `svds` returns $\tilde{U}_k, \tilde{\Sigma}_k, \tilde{V}_k$ for which $\|A\tilde{V}_k - \tilde{U}_k\tilde{\Sigma}_k\|_2 \leq \text{tol} \cdot \|A\|_2$. The accuracy and running time results are shown in Figure 1 which illustrates the linear speedup of our algorithm. For reference, at $n = 1000$, the exact solution takes about 2.5s on a single CPU laptop. Our algorithm performs according to the theory (with $p = O(\ln n)$ we achieve approximately fixed relative error).

References

- [1] Y. S. Abu-Mostafa, M. Magdon-Ismail, and H.-T. Lin. *Learning from data*. AMLBook, 2012.
- [2] C. Boutsidis and M. Magdon-Ismail. A note on sparse least-squares regression. *Information Processing Letters*, 2013.

- [3] T. F. Chan and P. C. Hansen. Some applications of the rank revealing QR factorization. *SIAM Journal on Scientific and Statistical Computing*, 13:727–741, 1992.
- [4] K. R. Davidson and S. J. Szarek. Local Operator Theory, Random Matrices and Banach Spaces. In *Handbook of the Geometry of Banach Spaces*, volume 1. Elsevier Science, 2001.
- [5] A. Gittens, P. Kambadur, and C. Boutsidis. Approximate spectral clustering via randomized sketching. *preprint arXiv:1311.2854*, 2013.
- [6] G. Golub and C. Van Loan. *Matrix computations*. JHU Press, 2012.
- [7] M. Gu and S. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM Journal on Scientific Computing*, 17:848–869, 1996.
- [8] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [9] P. Hansen. The truncated SVD as a method for regularization. *BIT Numerical Mathematics*, 27(4):534–553, 1987.
- [10] P. C. Hansen. Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numerical rank. *SIAM Journal on Scientific and Statistical Computing*, 11(3):503–518, 1990.
- [11] R. Horn and C. Johnson. *Matrix Analysis*. CUP, NY, 1985.
- [12] D. P. O’Leary and J. A. Simmons. A bidiagonalization-regularization procedure for large scale discretizations of ill-posed problems. *SIAM Journal on Scientific and Statistical Computing*, 2(4):474–489, 1981.
- [13] V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for principal component analysis. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1100–1124, 2009.
- [14] A. Sankar, D. A. Spielman, and S.-H. Teng. Smoothed analysis of the condition numbers and growth factors of matrices. *SIAM Journal on Matrix Analysis and Applications*, 28(2):446–476, 2006.
- [15] G. Stewart. On the perturbation of pseudo-inverses, projections and linear least squares problems. *SIAM review*, 19(4):634–662, 1977.
- [16] J. M. Varah. On the numerical solution of ill-conditioned linear systems with applications to ill-posed problems. *SIAM Journal on Numerical Analysis*, 10(2):257–267, 1973.
- [17] C. Vogel and J. Wade. Iterative SVD-based methods for ill-posed problems. *SIAM Journal on Scientific Computing*, 15(3):736–754, 1994.
- [18] H. Xiang and J. Zou. Regularization with randomized SVD for large-scale discrete inverse problems. *Inverse Problems*, 29(8):085008, 2013.

A Proof of Lemma 7

The result appeared in prior work [5, Corollary 11]. Nevertheless, for completeness, we give a short, different proof based on [6, Theorem 2.6.1], which states that for any two $m \times k$ orthonormal matrices \mathbf{W}, \mathbf{Z} with $m \geq k$:

$$\|\mathbf{W}\mathbf{W}^T - \mathbf{Z}\mathbf{Z}^T\|_2 = \|\mathbf{Z}^T\mathbf{W}^\perp\|_2 = \|\mathbf{W}^T\mathbf{Z}^\perp\|_2.$$

$\mathbf{Z}^\perp \in \mathbb{R}^{m \times (m-k)}$ is such that $[\mathbf{Z}, \mathbf{Z}^\perp] \in \mathbb{R}^{m \times m}$ is a full orthonormal basis. We set $\mathbf{U}_k^\perp = [\mathbf{U}_{\rho-k}, \mathbf{U}_{m-\rho-k}]$.

Given some (any) $\mathbf{S} \in \mathbb{R}^{n \times k}$, recall that in our algorithm, \mathbf{Q} is obtained by a QR-factorization of $(\mathbf{A}\mathbf{A}^T)^p \mathbf{A}\mathbf{S}$:

$$(\mathbf{A}\mathbf{A}^T)^p \mathbf{A}\mathbf{S} = \mathbf{Q}\mathbf{R}.$$

where $\mathbf{Q} \in \mathbb{R}^{m \times k}$ and $\mathbf{R} \in \mathbb{R}^{k \times k}$. We need some basic facts:

$$\mathbf{Q}\mathbf{R} = \mathbf{U}_k \Sigma_k^{2p+1} \mathbf{V}_k^T \mathbf{S} + \mathbf{U}_{\rho-k} \Sigma_{\rho-k}^{2p+1} \mathbf{V}_{\rho-k}^T \mathbf{S}; \quad (8)$$

$$\sigma_k(\mathbf{Q}\mathbf{R}) \geq \sigma_k(\mathbf{U}_k \Sigma_k^{2p+1} \mathbf{V}_k^T \mathbf{S}) \geq \sigma_k^{2p+1} \sigma_k(\mathbf{V}_k^T \mathbf{S}); \quad (9)$$

$$\sigma_i(\mathbf{Q}\mathbf{R}) = \sigma_i(\mathbf{R}); \quad (10)$$

$$\|\mathbf{X}\mathbf{R}\|_2 \geq \|\mathbf{X}\|_2 \sigma_k(\mathbf{R}), \quad \text{for any } \mathbf{X} \in \mathbb{R}^{\ell \times k}. \quad (11)$$

(8) follows from a direct computation using the SVD of \mathbf{A} ; (9) follows from (8) because \mathbf{U}_k and $\mathbf{U}_{\rho-k}$ span orthogonal spaces, and the fact that the minimum singular value of a product is at least the product of the minimum singular values; (10) follows because $\mathbf{Q}^T \mathbf{Q} = \mathbf{I}_k$; (11) is well known: it is clear if $\sigma_k(\mathbf{R}) = 0$ and if $\sigma_k(\mathbf{R}) > 0$ then it follows from:

$$\|\mathbf{X}_2\| = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{X}\mathbf{R}\mathbf{x}\|_2}{\|\mathbf{R}\mathbf{x}\|_2} \leq \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{X}\mathbf{R}\mathbf{x}\|_2}{\sigma_k(\mathbf{R}) \|\mathbf{x}\|_2} = \frac{\|\mathbf{X}\mathbf{R}\|_2}{\sigma_k(\mathbf{R})}.$$

Observe that $\tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T = \mathbf{Q} \mathbf{U}_{\mathbf{Q}^T \mathbf{A}} \mathbf{U}_{\mathbf{Q}^T \mathbf{A}}^T \mathbf{Q}^T = \mathbf{Q} \mathbf{Q}^T$, because $\mathbf{U}_{\mathbf{Q}^T \mathbf{A}} \mathbf{U}_{\mathbf{Q}^T \mathbf{A}}^T \mathbf{A} = \mathbf{I}_k$. Therefore, using [6, Theorem 2.6.1],

$$\begin{aligned} \|\mathbf{U}_k \mathbf{U}_k^T - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T\|_2 &= \|\mathbf{U}_k \mathbf{U}_k^T - \mathbf{Q} \mathbf{Q}^T\|_2 = \|\mathbf{Q}^T \mathbf{U}_k^\perp\|_2 \\ &= \|(\mathbf{U}_k^\perp)^T \mathbf{Q}\|_2 = \|\mathbf{U}_{\rho-k}^T \mathbf{Q}\|_2. \end{aligned} \quad (12)$$

The last equality is because $\mathbf{U}_{m-\rho-k}^T \mathbf{Q} = \mathbf{0}$ because \mathbf{Q} is in the range of \mathbf{A} . We now bound $\|\mathbf{U}_{\rho-k}^T \mathbf{Q}\|_2$.

$$\begin{aligned} \|\mathbf{U}_{\rho-k}^T \mathbf{Q}\mathbf{R}\|_2 &\geq \|\mathbf{U}_{\rho-k}^T \mathbf{Q}\|_2 \sigma_k(\mathbf{R}) \\ &\geq \|\mathbf{U}_{\rho-k}^T \mathbf{Q}\|_2 \sigma_k^{2p+1} \sigma_k(\mathbf{V}_k^T \mathbf{S}). \end{aligned} \quad (13)$$

$$\begin{aligned} \|\mathbf{U}_{\rho-k}^T \mathbf{Q}\mathbf{R}\|_2 &= \|\Sigma_{\rho-k}^{2p+1} \mathbf{V}_{\rho-k}^T \mathbf{S}\|_2 \\ &\leq \sigma_{k+1}^{2p+1} \sigma_1(\mathbf{V}_{\rho-k}^T \mathbf{S}). \end{aligned} \quad (14)$$

(13) follows using (11) then (10) then (9); (14) uses (8) and submultiplicativity. Using (12) with (13) and (14), we have:

Lemma 10. For any matrix $\mathbf{S} \in \mathbb{R}^{n \times k}$,

$$\sigma_k(\mathbf{V}_k^T \mathbf{S}) \|\mathbf{U}_k \mathbf{U}_k^T - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T\|_2 \leq \gamma_k^{2p+1} \sigma_1(\mathbf{V}_{\rho-k}^T \mathbf{S}).$$

Lemma 10 holds for general \mathbf{S} . We now use the fact that \mathbf{S} is a matrix of i.i.d. standard Gaussians. Then, for any orthonormal matrix \mathbf{V} , $\mathbf{V}^T \mathbf{S}$ is a matrix of i.i.d standard Gaussians. So, $\mathbf{V}_k^T \mathbf{S}$ is a $k \times k$ matrix to which Lemma 4 applies. Let $\mathbf{V} \in \mathbb{R}^{n \times n}$ be the extension of $\mathbf{V}_{\rho-k}$ to a full orthonormal basis. Then, $\mathbf{V}^T \mathbf{S}$ is an $n \times k$ matrix to which Lemma 3 applies (we set $t = 4$). By a union bound, with probability at least $1 - e^{-2n} - 2.35\delta$, both inequalities hold:

$$\begin{aligned}\sigma_k(\mathbf{V}_k^T \mathbf{S}) &\geq \delta k^{-\frac{1}{2}}; \\ \sigma_1(\mathbf{V}_{\rho-k}^T \mathbf{S}) &\leq \sigma_1(\mathbf{V}^T \mathbf{S}) \leq 4n^{1/2}\end{aligned}$$

Using Lemma 10 we conclude:

$$\|\mathbf{U}_k \mathbf{U}_k^T - \tilde{\mathbf{U}}_k \tilde{\mathbf{U}}_k^T\|_2 \leq 4\gamma_k^{2p+1} \delta^{-1} \sqrt{nk} \leq 4\gamma_k^{2p} \delta^{-1} n.$$

(We used $\gamma_k \leq 1$ and $k \leq n$.) Set $4\gamma_k^{2p} \delta^{-1} n = \varepsilon$ and solve for p to get $p = \ln(\varepsilon \delta / 4n) / \ln(\gamma_k^2)$, as desired.

■